

## DETAILED DESCRIPTION

[0013] Exemplary embodiments of the present invention relate generally to continual learning in slowly-varying environments, and more particularly, to providing an optimized action to an external system. Exemplary embodiments recognize without such a system to continually learn, human intervention and expertise is needed to monitor and adapt to changes in the environment. Exemplary embodiments for continual learning when the behavior of the environment varies slowly with time are described below with references to FIGS. 1-3.

[0014] Implementation of such exemplary embodiments may take a variety of forms, and exemplary implementation details are discussed subsequently with reference to the Figures.

[0015] FIG. 1 is a functional block diagram illustrating a data processing environment 100, according to an exemplary embodiment. FIG. 1 provides an illustration of one implementation of the subject matter and does not imply any limitations in which different exemplary embodiments of the data processing environment may be implemented. Many modifications to the depicted data processing environment may be made by those skilled in the art without departing from the scope and range of equivalents of the subject matter. In some exemplary embodiments, data processing environment 100 includes a network 106, a server 104, which operates continual learning system 102, and a decision agent 108.

[0016] Network 106 interconnects server 104 and decision agent 108. In general, network 106 can be any combination of connections and protocols capable of supporting communications between server 104, decision agent 108, and continual learning system 102. Network 106 can include wire cables, wireless communication links, fiber optic cables, routers, switches, firewalls, or any combination that can include wired, wireless, or fiber optic connections known by those skilled in the art.

[0017] Server 104 can be a web-based server hosting continual learning system 102, in accordance with exemplary embodiments of the present invention. In one exemplary embodiment, server 104 can be any programmable electronic device or computing system capable of receiving and sending data, via network 106, and performing computer-readable program instructions known by those skilled in the art. In some exemplary embodiments, server 104 can include a data storage repository (not shown) for storing data including, but not limited to, state information for all entities associated with an environment, transaction data, and various models or policies. Data storage repository can be any programmable electronic device or computing system capable of receiving, storing, and sending files and data, and performing computer readable program instructions capable of communicating with server 104 and decision agent 108, via network 106. In some exemplary embodiments, server 104 can be a cluster of computer nodes, in the distributed system, operating continual learning system 102, via network 106. In an exemplary embodiment, server 104 includes databases (not shown) that provides a service to external systems. In another exemplary embodiment, continual learning system 102 resides locally on server 104. In yet another exemplary embodiment, continual learning system 102 resides locally on one or more computer nodes.

[0018] In some exemplary embodiments, server 104 includes continual learning system 102 that utilizes compo-

nents or models in order to select and provide an optimized action to decision agent 108. For example, continual learning system 102 utilizes a decision engine 202 to perform various functions for receiving one or more action requests from a decision agent 108. Continual learning system 102, utilizing the decision engine 202, can select a model providing an optimized action. Continual learning system 102, utilizing orchestrator 212, can deploy a first model to the decision engine 202, initiate an observation period, and initiate a test period. In another example, continual learning system 102 utilizes a model builder 210 to build the second model. In a final example, continual learning system 102 utilizes an evaluator 214 to determine a performance score for the first model and a performance score for the second model.

[0019] In some exemplary embodiments, continual learning system 102 operates on a central server, such as server 104, and can be utilized by one or more decision agents, such as decision agent 108, via a mobile application downloaded from the central server or a third-party application store, and executed on the one or more decision devices. In another exemplary embodiment, continual learning system 102 may be a software-based program, downloaded from a central server, such as server 104, and installed on one or more decision agents, such as decision agent 108. In yet another exemplary embodiment, continual learning system 102 can be utilized as a software service provided by a third-party cloud service provider (not shown).

[0020] In some exemplary embodiments, decision agent 108 is an agent to the server 104 and can be for example, a desktop computer, a laptop computer, a smart phone, or any other electronic device or computing system, known by those skilled in the art, capable of communicating with the server 104 through the network 106. For example, decision agent 108 may be a laptop computer capable of connecting to a network, such as network 106, to access continual learning system 102 and to provide requests for actions and rewards. In other exemplary embodiments, decision agent 108 can be any suitable types of mobile devices capable of running mobile applications or a mobile operating system.

[0021] In an exemplary embodiment, decision agent 108 interacts with an environment (e.g. a clothing retailer website or a news website). Decision agent 108 may request continual learning system 102 for an action. In another exemplary embodiment, decision agent 108 may provide the state of the environment (e.g. the duration of a user on articles of the news website, or the dollar amount of clothing purchases by a user) in the request. It can also be appreciated by those skilled in the art that decision agent 108 can provide an identifier of the entity corresponding to the request in lieu of the state of the environment.

[0022] FIG. 2 is a functional block diagram depicting components of continual learning system 102, according to an exemplary embodiment. In some exemplary embodiments, the components reside locally with continual learning system 102. In another exemplary embodiments, the components, either individually or collectively, reside remotely from continual learning system 102 and are connected to continual learning system 102, via network 106.

[0023] In some exemplary embodiments, Decision Engine 202 receives a request for an action from Decision Agent 108. Decision Engine 202 retrieves parameters associated with a model from Model Database 208. Model Builder 210 builds a first model offline which it seeds to Model Database